

There are a number of mitigation tactics available. At the very least, application providers may choose to develop user documentation or customer service expertise to help their users understand the nature of the problem and potential workarounds, if available. Another tactic may be to redesign some applications to use different ports, to conduct connectivity testing before establishing connections, to be port-agile, or to make port selection user-configurable. Whether any of these options are available may depend on whether re-designed versions of the application can be made compatible with existing versions.

These mitigations may raise additional issues for application providers. Introducing connectivity checks can impact performance, causing applications to take a significantly longer amount of time to establish initial connections. Shifting to ports that are already in common use by other applications and protocols can complicate application design. For example, some networks use proxies to validate that HTTP traffic conforms to specific protocol semantics; shifting non-HTTP traffic to port 80 may therefore result in the loss of particular functionality or may prevent the use of non-TCP transports. The implications of movement towards the majority of applications running on a small number of ports are uncertain as of yet. Such a change could arguably dampen the “diversity” or limit the number of different types of applications that can perform well on the network, since new applications may be expected to conform to the way that existing applications function on the same port (for example, expecting that all TCP /80 traffic behaves like HTTP).

Port blocking on residential networks may particularly constrain independent or non-commercial application developers, many of whom experiment with new application features and functionality using residential broadband connections. Although the rise of cloud computing resources may provide these developers with a way to circumvent restrictions imposed on their home connections, port blocking on residential networks may still put limits on local testing and development.

Port blocking is among a set of tools and tactics (NAT being the other major example) that can undermine the original intent of ports: to provide reliable local addresses so that end systems could manage multiple communications at once. In general, blocking ports does not cause applications to vanish from the Internet, but rather induces a cat-and-mouse game whereby application development either becomes increasingly complex so as to evade port blocking through port-agnosticism, or drives application traffic to a dwindling set of ports that are reliably kept open across most networks. These efforts in turn cause ISPs to seek increasingly application-aware means of identifying and thwarting unwanted traffic.

4.1.3. Consumer or End-User Concerns

Port blocking can cause applications to not function properly, or “break”, by preventing them from using the ports they were designed to use. Importantly, it may not be obvious to Internet users why their affected application is not working because

the application may simply be unable to connect or fail silently. If error messages are provided, those messages may not contain specific details about the cause of the problem. Users may seek assistance from the ISP's customer service, online documentation, or other knowledgeable sources if they cannot diagnose the problem themselves. The process of diagnosis is further complicated by the fact that the problem could alternatively be caused by home networking equipment or a software-based port block.

Users' ability to respond to port blocking depends on their technical sophistication and the extent to which workarounds are available. Overcoming the port block may require installing a software update, changing a configuration setting, requesting an opt-out from the ISP, or upgrading the level of service (from residential to business, for example). If these options are not available, or if users lack the knowledge or willingness to pursue them, they may be prevented from using the blocked application altogether, or they may have to switch to a different application or a different network (from wired to wireless, for example). Where port blocking is used to funnel traffic to an ISP's own infrastructure (by limiting outbound TCP /25 traffic unless the traffic is destined for the ISP's own mail servers, for example), it effectively reduces the set of application provider choices available to users (all other mail servers, for example).

The trend towards port overloading, or in other words the fact that many different applications now use the same port, means that traffic identification and classification need to take place at the application layer. This may have implications for user control and privacy, because port overloading motivates the deployment of Deep Packet Inspection (DPI) and other content-aware technologies that can be used to identify and manage specific applications or communications.

Blocking of certain ports could also have more serious repercussions for user privacy and security. For example, blocking port 443 would effectively prevent secure HTTP communication and the ability of users to connect with the large number of sites that require Hypertext Transport Protocol Secure (HTTPS). Port blocking used in this capacity is an attempt to keep communications in clear text (perhaps for inspection or surveillance purposes). This port has been blocked by ISPs outside of the US, but not domestically.

5. Technical Working Group (TWG) Suggested Practices

While port blocking can have positive security benefits, it can affect how particular Internet applications function. Thus its use has the potential to be anti-competitive, discriminatory, otherwise motivated by non-technical factors, or construed as such. As a result, the Broadband Internet Technical Advisory Group (BITAG) has a number of suggested practices regarding port blocking on both wireline and wireless networks.

5.1. ISPs Should Avoid Port Blocking Unless No Reasonable Alternatives Are Available

BITAG recommends that ISPs avoid port blocking unless they have no reasonable alternatives available for preventing unwanted traffic and protecting customers. Further, if port blocking is deemed necessary, it should only be used for the purposes of protecting the implementing ISP's network and users. Port blocking should not be used for ongoing capacity management, or to enforce non-security terms of service, or to disadvantage competing applications.

Port blocking can create collateral damage for legitimate users and uses of the network, and can complicate the development of applications. A number of applications (including many that pose security threats) have evolved to become port-agile or to use ports that are unlikely to be blocked, most commonly ports 80 and 443. As a result, the long-term effectiveness of port blocking as a means to prevent unwanted traffic is limited. On the other hand, ISPs may view port blocking as a simple and powerful way of handling security threats, particularly in the short term. Despite the negative impacts that may come with the practice, for the time being its use may be considered a necessity.

5.2. ISPs Should Provide Opt-Out Provisions

BITAG recommends that if an ISP can reasonably provide their users with opt-out provisions or exceptions to their port blocking policies, they should do so.

BITAG recognizes the benefit of providing opt out policies for a subset of users and for certain port blocking rules. However, the technical feasibility, administrative complexity, and costs can vary greatly depending upon the implementation, including the particulars of the access network technology (i.e. DOCSIS, DSL, LTE, etc.). For example, there may be cases where SNMP blocking is only feasible in an access router or other aggregation point, some distance from the user's equipment, which may make per-user controls exceedingly difficult or impossible. Thus, BITAG recommends that ISPs balance their decisions about which ports to block with their capabilities of offering opt-out.

5.3. ISPs Should Disclose Port Blocking Policies

BITAG recommends that ISPs publicly disclose their port blocking policies. The information should be readily available to both customers and non-customers alike, and should be as informative and concise as possible. For example, port blocking policies could be provided on the ISP's public facing website, on a page dedicated to summarizing or describing the respective ISP's network management practices.

For persistent port blocks the information should include:

- Port number(s)
- Transport protocol (e.g., TCP or UDP)
- Application(s) normally associated with the port(s) (e.g., SMTP)
- Direction of the block (outbound or inbound)
- Brief description of the reason(s) for the block (e.g., SPAM)
- If opt-out provisions are available and how to request such

This will give users better information with which to diagnose problems, can better inform consumers' decision-making when choosing an ISP, and can provide crucial information to application developers.

There may be times when a security incident will prompt the need for an immediate and temporary port block to be implemented by an ISP in order to protect its customers or protect its network. It may not be feasible to disclose such blocks before they are removed. There may also be times when the disclosure of port blocking in response to a particular attack may compromise that security mitigation.

5.4. ISPs Should Make Communications Channels Available for Feedback

BITAG recommends that ISPs provide a communications channel or other clear method for application providers and consumers to provide feedback to each ISP on its respective port blocking policy – to discuss impacts caused by port blocking and to consider other possible mitigations, among other things. The communications channel or other clear methods should be provided where the port blocking policies are disclosed. ISPs should be reasonably responsive to communications received from application developers and consumers, among other things to discuss impacts caused by port blocking and to consider possible mitigations.

5.5. ISPs Should Revisit Their Port Blocking Policies on a Regular Basis

BITAG recommends that ISPs revisit their respective port blocking policies on a regular basis to determine whether the threats that required the port blocking rules continue to be relevant, and whether their policies should be adjusted accordingly. Some security threats are permanent and some are transitory or short-lived. Items such as spam prevention by blocking TCP/25 are expected to last quite some time, while others such as blocks to prevent certain types of malware may be temporary and can be fixed over time with software patching.

5.6. Port Blocking Rules for Consumer Equipment Should Be User Configurable

BITAG recommends that the port blocking (or firewall) rules of consumers' home routers should be user configurable – whether the routers are provided by the ISP or purchased separately by the consumer. It is recommended that the documentation provided with each unit inform the consumer that port blocking or firewall rules have been implemented, default ports blocked, and how consumers can modify those rules.

6. References

- [BCP38] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, May 2000, <<http://tools.ietf.org/html/bcp38>>.
- [BCP165] Cotton, M., L. Eggert, J. Touch, M. Westerlund, and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP165, August 2011, <<http://tools.ietf.org/html/bcp165>>.
- [RFC675] Cerf, V., Y. Dalal, and C. Sunshine, "Specification of Internet Control Program", RFC 675, December 1974, <<http://tools.ietf.org/html/rfc675>>.
- [RFC768] Postel, J., "User Diagram Protocol", RFC 768, August 1980, <<http://tools.ietf.org/html/rfc768>>.
- [RFC788] Postel, J., "Simple Mail Transfer Protocol", RFC 788, November 1981, <<http://tools.ietf.org/html/rfc788>>.
- [RFC793] Postel, J., "Transmission Control Protocol", RFC 793, September 1981, <<http://www.ietf.org/rfc/rfc793.txt>>.
- [RFC1001] Aggarwal, A., et al., "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods", RFC 1001, March 1987, <<http://www.ietf.org/rfc/rfc1001.txt>>.
- [RFC1002] Aggarwal, A., et al., "Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications", RFC 1002, March 1987, <<http://www.ietf.org/rfc/rfc1002.txt>>.
- [RFC1122] Braden, R., "Requirements for Internet Hosts – Communications Layers", RFC 1122, October 1989, <<http://tools.ietf.org/html/rfc1122>>.
- [RFC2616] Fielding, R., J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext Transfer Protocol – HTTP/1.1", RFC 2616, June 1999, <<http://tools.ietf.org/html/rfc2616>>.
- [RFC2827] Ferguson, P., D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing", RFC 2827, May 2000, <<https://tools.ietf.org/html/rfc2827>>.
- [RFC4271] Rekhter, Y., T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006, <<http://tools.ietf.org/html/rfc4271>>.
- [RFC5068] Hutzler, C., D. Crocker, P. Resnick, E. Allman, T. Finch, "Email Submission

- Operations: Access and Accountability Requirements", RFC 5068, November 2007, <<https://tools.ietf.org/html/rfc5068>>.
- [RFC6204] Singh, H., W. Beebe, C. Donley, B. Stark, and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011, <<http://tools.ietf.org/html/rfc6204>>.
- [RFC6335] Cotton, M., et al., "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", RFC 6335, August 2011, <<http://tools.ietf.org/html/rfc6335>>.
- [RFC6409] Gellens, R., and J. Klensin, "Message Submission for Mail", RFC 6409, November 2011, <<http://tools.ietf.org/html/rfc6409>>.
- [BITAG Large Scale NAT Report] Broadband Internet Technical Advisory Group (BITAG), "Implications of Network Address Translation (NAT)", March 2012, <http://www.bitag.org/documents/BITAG_TWG_Report-Large_Scale_NAT.pdf>.
- [BITAG SNMP Report] Broadband Internet Technical Advisory Group (BITAG), "SNMP Reflected Amplification DDoS Attack Mitigation", March 2012, <<http://www.bitag.org/documents/SNMP-Reflected-Amplification-DDoS-Attack-Mitigation.pdf>>.
- [Comcast Letter on SMTP Port 25] O'Reirdan, M., "Updated Management of SMTP Port 25", August 1, 2012, <<http://corporate.comcast.com/comcast-voices/updated-management-of-smtp-port-25>>.
- [M3AAWG Port 25 Recommendation] Messaging Malware Mobile Anti-Abuse Working Group (M3AAWG), "Managing Port 25 for Residential or Dynamic IP Space Benefits of Adoption and Risks of Inaction", December 2005, <http://www.maawg.org/sites/maawg/files/news/MAAWG_Port25rec0511.pdf>.
- [Madison River] Federal Communications Commission (FCC), "In the Matter of Madison River Communications, LLC and Affiliated Companies", Consent Decree, DA 05-543, March 2005, <http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-05-543A2.pdf>.
- [Netalyzr] University of California – Berkeley, International Computer Science Institute, Netalyzer", <<http://netalyzr.icsi.berkeley.edu/>>.
- [Netalyzr2010] Kreibich, C., N. Weaver, B. Nechaev, V. Paxson, "Netalyzer: Illuminating the Edge Network", November 2010, <<http://www.icir.org/christian/publications/2010-imc-netalyzr.pdf>>.
- [Port Number Registry] Touch, J., M. Kojo, E. Lear, A. Mankin, K. Ono, M. Stiernerling, and L.

Eggert, "Service Name and Transport Protocol Port Number Registry", March 2013, <<http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>>.

[Skype FAQ] Skype, "Connection Problems: Which Ports Need to be Open to Use Skype for Windows Desktop", April 2013, <<https://support.skype.com/en/faq/FA148/which-ports-need-to-be-open-to-use-skype-for-windows-desktop>>.

[Toward Quantifying Network Neutrality] Beverly, R., S. Bauer, A. Berger, "The Internet's Not a Big Truck: Toward Quantifying Network Neutrality", 2007, <http://www.akamai.com/dl/technical_publications/truck-pam07.pdf>.

[SANS] SANS Institute, "Intrusion Detection FAQ: What Port Numbers Do Well-Known Trojan Horses Use?", April 2013, <<http://www.sans.org/security-resources/idfaq/oddports.php>>.

7. Glossary of Terms

- **Home Gateway Device:** A network element that creates, connects to, or extends a home network for a user. These devices can perform a range of functions, such as connecting to the Internet, creating or extending a wireless network, providing backup and storage, etc. [See also RFC 6204]
- **HTTP Proxy:** A computer system or an application that acts as an intermediary for requests from clients seeking resources from other servers. A client connects to the proxy server, requesting some service, such as a file, connection, web page, or other resource available from a different server and the proxy server evaluates the request as a way to simplify and control its complexity.
- **ISP Interconnection Links:** For the purpose of this document, the places (links) where IP traffic is exchanged between ISP networks.
- **Transmission Control Protocol (TCP):** A protocol used along with the Internet Protocol (IP) to send data in the form of information packets between computers over the Internet. While IP handles the actual delivery of the data, TCP keeps track of the individual packets that a message is divided into for efficient routing through the Internet. IP packets can be lost, duplicated, or delivered out of order and TCP detects these problems, requests retransmission of lost data, rearranges out-of-order data, and even helps minimize network congestion to reduce the occurrence of the other problems. [See also RFC 675 et al]
- **User Datagram Protocol (UDP):** A protocol used along with the Internet Protocol (IP) to send data in the form of information packets between computers over the Internet. In contrast to TCP, UDP uses a simple transmission model with a minimum

of protocol mechanism. UDP is suitable for purposes where error checking and correction is either not necessary or performed in the application, thus avoiding the overhead of such processing at the network interface level. Time-sensitive applications often use UDP, where dropping packets is preferable to waiting for delayed packets. [See also RFC 768]

8. Document Contributors and Reviewers

- Fred Baker, Cisco
- Alissa Cooper, Center for Democracy and Technology
- Chuck Dvorak, AT&T
- Michael Fargano, CenturyLink
- David Fullagar, Netflix
- Jeffrey Good, Disney
- Amer Hassan, Microsoft
- Dale Hatfield
- Trace Hollifield, Bright House Networks
- Scott Jordan, University of California, Irvine
- Kevin Kahn, Intel
- Jason Livingood, Comcast
- Donald Smith, CenturyLink
- Jeff Swinton, Verizon
- Tony Watson, Google
- Jason Weil, Time Warner Cable